Rule WLM172: Server was idle a significant percent of time

Finding: The service class period identified in Rule WLM104 or WLM105 missed its response goal. However, address spaces handled by the server service class were Idle for a significant percent of their overall active time. Consequently, the Workload Manager delay information for the server service class may be meaningless.

Impact: This finding is provided for information purposes.

Logic flow: The following rules cause this rule to be invoked:

Rule WLM101:	Service Class did not achieve average response goal
Rule WLM102:	Service Class did not achieve percentile response goal

Discussion: When CPExpert produces Rule WLM104 or Rule WLM105 to indicate that a transaction service class did not achieve its response performance goal, the logic of these rules tries to identify the cause of the delay. The cause of the delay initially is analyzed from the "served" service class view. The delays from the served service class are reported by CICS (with CICS/ESA Version 4.1 or later, IMS (with IMS Version 5 or later), or DB2 Version 6 or later. The reporting is done by interaction with the Workload Manager, using the Workload Management Services macros[1].

Please refer to Rule WLM120 to Rule WLM132 for a discussion of the delays from the served service class.

After analyzing the **served** service class delays, CPExpert identifies the **server** service class. The server service class normally will be one or more CICS regions or IMS regions. The subsystem service class (e.g., the CICS region or IMS region) must have a performance goal and importance defined, in order for the region to start-up. However, the performance goal and importance normally are used by the Workload Manager **only at start-up time** for the address space[2].

After start-up time, the Workload Manger normally ignores the goal and importance of subsystems. After start-up time, the Workload Manager

[1]Please refer to Section 4 of this document for more detail about the Workload Management Services macros and how the subsystems use these macros to exchange information with the Workload Manager.

[2]This statement is not true if the region should become idle for some period of time. If there are no transactions executing in the region for some time, the Workload Manager will rely on the performance goal and importance associated with the region to make resource allocation decisions. This situation should normally occur only during "off shifts" or for test regions with low activity.

normally uses the goal and importance of the "served" transaction service classes as the basis for it resource allocation decisions.

The Workload Manager attempts to meet the performance goals of the "served" transaction service classes. In order to meet these performance goals, the Workload Manager must assign resources to the server service class (e.g., the service class of the CICS region), regardless of the goal and importance assigned to the subsystem service class.

The Workload Manager periodically examines the SRM control blocks describing each address space and acquires samples[3] describing the state of each dispatchable unit of an address space (that is, each TCB or SRB associated with the address space). The Workload Manager accumulates the samples into counters that describe the state of the address space. The samples are summarized by service class period.

The analysis performed by the Workload Manager and subsequent analysis by CPExpert is based on samples. The reliability of sampling depends upon having a sufficiently large number of samples such that the samples represent the "population" being sampled[4]. If a small number of samples are taken, invalid conclusions might be reached based on an analysis of the samples. In order for the conclusions about causes of delays to be valid, sufficient samples must be taken while address spaces were in a "ready" state rather than in an "idle" state.

When CPExpert determines that a transaction service class has missed its performance goal, CPExpert reviews the number of samples taken during times when address spaces in the **server** were in a "ready" state.This number of samples is obtained by summing the CPU Using samples (R723CCUS), I/O Using samples (R723CIOU), non-DASD I/O Using or Delay samples (R723CNDI), Total Delay samples (R723CTOT), and Unknown samples (R723CUNK). CPExpert produces Rule WLM172 if this total number of samples is small.

Once CPExpert has determined that an unacceptably small number of samples exist, no further analysis is done. It makes no sense to analyze delays to the service class based on a low number of samples, inasmuch as the conclusions from the samples would be invalid.

The following example illustrates the output from Rule WLM172:

[3]With MVS/ESA SP5.1 Goal Mode, the sampling is done every 250 milliseconds. The sampling interval is recorded in SMF Type 72 records (R723MTVL).

[4]With the Workload Manager samples, the "population" consists of the possible execution states of address spaces being sampled.


```
RULE WLM172:  SERVER WAS IDLE A SIGNIFICANT PERCENT OF TIME

  The delay information presented above is based on the EXECUTION time of
  the CICSTEST server (the CPU Using, Execution Delay, and Unknown Delay).
  These percentages show the distribution of time while some transaction
  was active.  However, address spaces in the CICSTEST Service Class
  were IDLE for a significant percent of their overall active time.  The
  below information shows the percent of CICSTEST total active time in
  which address spaces were executing (processing transactions) or were
  idle.  Please refer to Rule WLM172 in the WLM Component User Manual for
  a discussion of the implications of this finding.

                          AVERAGE      PCT         PCT       PCT    EXECUTION
   MEASUREMENT INTERVAL     MPL     EXECUTING   UNKNOWN    IDLE     SAMPLES
   11:15-11:29,07DEC1994      1         0.1         .5      99.9         6
```

Suggestion: CPExpert suggests that you consider the following alternatives:

- You can ignore the finding (and previous rules in the logic flow) if you feel that the situation is unusual rather than a continuing status. For example, the finding might be made when a server was temporarily idle because development personnel were not submitting transactions to the CICS test region.

  If you chose to ignore the finding, you may wish to exclude the transaction service class from analysis, using the EXCLUDE guidance parameters described in Section 2 (Chapter 1.1.8) of this document). You likely would become annoyed by CPExpert continually reporting that the service class missed its performance goal when you contemplate no action.

- You may wish to delete the service class and assign the workload to a service class with more active address spaces if you feel that the situation is a continuing one. That is, if you feel that the address spaces normally are idle, you may wish to review whether they need their own service class. As general guidance, it is desirable to keep the service class periods to as small a number as possible.

